arXiv:1608.06863v1 [cs.CV] 24 Aug 2016

Kullback-Leibler Penalized Sparse Discriminant Analysis for Event-Related Potential Classification

Victoria Peterson*[1], Hugo Leonardo Rufiner[1,2], and Ruben Daniel Spies[3]

[1]Instituto de Investigación en Señales, Sistemas e Inteligencia Computacional, UNL, CONICET, FICH, Ruta Nac. 168, km 472.4, 3000, Santa Fe, Argentina.
[2]Facultad de Ingeniería, Universidad Nacional de Entre Ríos, Ruta Prov. 11, km 10, 3100, Oro Verde, Argentina.
[3]Instituto de Matemática Aplicada del Litoral, UNL, CONICET, FIQ, Predio Dr. Alberto Cassano del CCT-CONICET-Santa Fe, Ruta Nac. 168, km 0, 3000, Santa Fe, Argentina.

August 25, 2016

Abstract

A brain computer interface (BCI) is a system which provides direct communication between the mind of a person and the outside world by using only brain activity (EEG). The event-related potential (ERP)-based BCI problem consists of a binary pattern recognition. Linear discriminant analysis (LDA) is widely used to solve this type of classification problems, but it fails when the number of features is large relative to the number of observations. In this work

*vpeterson@sinc.unl.edu.ar

we propose a penalized version of the sparse discriminant analysis (SDA), called Kullback-Leibler penalized sparse discriminant analysis (KLSDA). This method inherits both the discriminative feature selection and classification properties of SDA and it also improves SDA performance through the addition of Kullback-Leibler class discrepancy information. The KLSDA method is design to automatically select the optimal regularization parameters. Numerical experiments with two real ERP-EEG datasets show that this new method outperforms standard SDA.

1 Introduction

A brain computer interface (BCI) is a system that measures brain activity and converts it into an artificial output which is able to replace, restore or improve any normal output (neuromuscular or hormonal) used by a person to communicate and control his/her external or internal environment. Thus, BCI can significantly improve the quality of life of people with severe neuromuscular disabilities [35].

Communication between the brain of a person and the outside world can be appropriately established by means of a BCI system based on event-related potentials (ERPs), which are manifestations of neural activity as a consequence of certain infrequent or relevant stimuli. The main reason for using ERP-based BCI are: it is non-invasive, it requires minimal user training and it is quite robust (in the sense that it can be use by more than 90 % of people) [34]. One of the main components of such ERPs is the P300 wave, which is a positive deflection occurring in the scalp-recorded EEG approximately 300 ms after the stimulus has been applied. The P300 wave is unconsciously generated and its latency and amplitude vary between different EEG records of the same person, and even more, between EEG records of different persons [18]. By using the "oddball" paradigm [8] an ERP-based BCI can decode desired commands from the subject by detecting those ERPs in the background EEG. In other words, a ERP-based BCI deals with a pattern recognition problem where two classes are involved: EEG with ERP (target class) and EEG without ERP (non-target class).

Several authors have investigated in regards to which classification technique performs best for separating target from non-target EEG records [25, 4, 24, 11]. All of these studies have concluded that linear discriminant anal-

ysis (LDA) is a very good classification scheme, resulting most of the times in the optimal performance while keeping the solution simple. It is timely to mention here that in a BCI classification scheme there are two main problems: the curse-of-dimensionality and the bias-variance trade-off [25]. While the former is a consequence of working with a concatenation of multiple time points from multiple channels, the latter depicts the generalization capability of the classifier.

However, it has been shown that LDA performs poorly when high dimensional data with small training sample is used. Different regularized solutions have been proposed [21, 19, 37] and all of them have demonstrated that a regularized version of LDA can substantially improve the classification performance of the standard LDA. However, there is still much room for further improvement. In this regards, in the present work, we propose to achieve significant gain by building up a model which takes into account the following issues, which, to our knowledge, have not been jointly considered in any of the previous works:

- Solution sparsity needs to be kept as high as possible in order to maximize robustness while achieving good classification performance. It is timely to remember that the generalizability of a model tends to decrease as the sparsity decreases, i.e., as the number of parameters in the model increases [17].

- Since randomly selecting one variable out of a group of correlated variables may produce neglecting important information, groups of "similar" variables should be allowed into the model [38].

- Model should allow for the inclusion of discriminant a priori information since that could greatly enhance separability between classes.

With the above in mind, in this work we develop a penalized version of the sparse discriminant analysis (SDA) [6], which we call Kullback-Leibler Penalized sparse discriminant analysis (KLSDA), with the main objective of solving the binary ERP classification problem. As far as we know SDA has never been used before in ERP-based BCI classification problems. Also, we shall compare our penalized version with the standard SDA in this context.

The organization of this article is as follows. In Section 2 we made a brief review on discriminant analysis. Our proposed new approach is presented in its full general formulation in Section 3. In Section 4 the ERP-EEG databases

used in the experiments are described. In Section 5 the experiments and results of the ERP-based EEG classification problem solved by KLSDA are shown. Finally, concluding remarks and future works are presented in Section 6.

2 Discriminant Analysis: a brief review

The LDA criterion is a well-known dimensionality reduction tool in the context of supervised classification. Its popularity is mainly due to its simplicity and robustness which lead to very high classification performances in many applications [7].

Let $\mathbf{W}_1, \ldots, \mathbf{W}_\mathrm{K}$ be p-dimensional random vectors whose distributions uniquely characterize each one of the K classes of a given classification problem. In addition, let $\mathbf{X}$ be an $n \times p$ data matrix such that each row $\mathbf{x}_i$ is a realization of one and only one of the aforementioned random vectors, and let $\mathbf{z} \in \mathbb{R}^n$ be a categorical variable accounting for class membership; i.e., such that if pattern $\mathbf{x}_i$ is a realization of $\mathbf{W}_\mathrm{k}$, then $z_i = k$.

The LDA method consists of finding $q < K$ discriminant vectors (directions), $\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_q$ such that by projecting the data matrix $\mathbf{X}$ over those directions, the "classes" will be well separated one from each other. It is assumed that the random vectors $\mathbf{W}_1, \ldots, \mathbf{W}_\mathrm{K}$ are independently and normally distributed with a common covariance matrix $\boldsymbol{\Sigma}_t$. The procedure for finding the vectors $\boldsymbol{\beta}_j$ requires estimates of the within-class, the between-class and the total covariance matrices, $\boldsymbol{\Sigma}_w$, $\boldsymbol{\Sigma}_b$ and $\boldsymbol{\Sigma}_t$, respectively. These estimates are given by:

$$\hat{\boldsymbol{\Sigma}}_w = \frac{1}{n} \sum_{k=1}^{K} \sum_{i \in I_k} (\mathbf{x}_i - \boldsymbol{\mu}_k)(\mathbf{x}_i - \boldsymbol{\mu}_k)^T,$$

$$\hat{\boldsymbol{\Sigma}}_b = \frac{1}{n} \sum_{k=1}^{K} n_k (\boldsymbol{\mu}_k - \boldsymbol{\mu})(\boldsymbol{\mu}_k - \boldsymbol{\mu})^T,$$

$$\hat{\boldsymbol{\Sigma}}_t = \frac{1}{n} \sum_{i=1}^{n} (\mathbf{x}_i - \boldsymbol{\mu})(\mathbf{x}_i - \boldsymbol{\mu})^T,$$

where I_k and n_k are the set of indices and the number of patterns belonging to class k, respectively, $\boldsymbol{\mu}_k \doteq \frac{1}{n_k} \sum_{i \in I_k} \mathbf{x}_i$ is the class k sample mean and $\boldsymbol{\mu} \doteq \frac{1}{n} \sum_{k=1}^{K} \boldsymbol{\mu}_k$ is the common sample mean. Note that $\hat{\boldsymbol{\Sigma}}_t = \hat{\boldsymbol{\Sigma}}_w + \hat{\boldsymbol{\Sigma}}_b$.

Let $\mathbf{B}$ be a $p \times q$ matrix whose j^{th} columns is $\boldsymbol{\beta}_j$. The LDA method seeks to find the vectors $\boldsymbol{\beta}_j$ in such a way that they maximize separability between classes, which is achieved by simultaneously maximizing $\hat{\boldsymbol{\Sigma}}_b$ and minimizing $\hat{\boldsymbol{\Sigma}}_w$. Since $\hat{\boldsymbol{\Sigma}}_t = \hat{\boldsymbol{\Sigma}}_w + \hat{\boldsymbol{\Sigma}}_b$, this is equivalent to simultaneously maximizing $\hat{\boldsymbol{\Sigma}}_b$ and minimizing $\hat{\boldsymbol{\Sigma}}_t$, thus:

$$\mathbf{B}^* = \underset{\mathbf{B} \in \mathbb{R}^{p \times q}}{\arg\max} (\mathbf{B}^T \hat{\boldsymbol{\Sigma}}_b \mathbf{B})(\mathbf{B}^T \hat{\boldsymbol{\Sigma}}_t \mathbf{B})^{-1}. \tag{1}$$

Since the rank of $\hat{\boldsymbol{\Sigma}}_b$ is at most $K-1$, there are at most $K-1$ non-trivial solutions of problem (1), and therefore there are at most $K-1$ discriminative vectors. Observe that, these vectors are precisely the directions along which the classes show maximum between-class covariance relative to their within-class covariance. Usually $q = K-1$.

It is known that LDA fails when the number of patterns is low in relation to the number of variables ($n << p$) [15]. In this situation, the matrix $\hat{\boldsymbol{\Sigma}}_t$ becomes usually ill-conditioned, what produces a poor estimation of $\mathbf{B}$ and, as a consequence, a bad classification performance.

In the particular case $K = 2$ (and therefore $q = 1$), the solution to (1) has the following explicit formulation:

$$\boldsymbol{\beta}^* = \hat{\boldsymbol{\Sigma}}_t^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2). \tag{2}$$

This special case is known as Fisher linear discriminant analysis (FLDA) [10]. The FLDA approach can be formulated as a linear regression model [10, 7]. Let $\mathbf{X}$ be as before and let $\mathbf{y}$ be a n-dimensional vector such that $y_i = \frac{n_2}{n}$ or $y_i = -\frac{n_1}{n}$, depending on whether the i^{th} observation belonging to class 1 or to class 2, respectively, and let us consider the following ordinary least squares problem (OLS):

$$\boldsymbol{\alpha}^* = \underset{\boldsymbol{\alpha} \in \mathbb{R}^p}{\arg\min} \|\mathbf{y} - \boldsymbol{X}\boldsymbol{\alpha}\|_2^2, \tag{3}$$

whose solutions are all the vectors in the set $\mathcal{N}(\mathbf{X}^T\mathbf{X}) + (\mathbf{X}^T\mathbf{X})^\dagger \mathbf{X}^T\mathbf{y}$, where "†" denotes the Moore-Penrose generalized inverse and $\mathcal{N}(\mathbf{X}^T\mathbf{X})$ denotes the null space of $\mathbf{X}^T\mathbf{X}$. If $\mathbf{X}^T\mathbf{X}$ is invertible, then Eq. (3) has a unique solution given by $\boldsymbol{\alpha}^* = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{y}$. For convenience it is assumed that $\boldsymbol{\mu} = 0$, and therefore $\mathbf{X}^T\mathbf{X} = n\boldsymbol{\Sigma}_t$ and $\mathbf{X}^T\mathbf{y} = \frac{n_1 n_2}{n}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$. Hence $\boldsymbol{\alpha}^* = \frac{n_1 n_2}{n^2}\boldsymbol{\beta}^*$, where $\boldsymbol{\beta}^*$ is given by (2). Since the proportionality constant $\frac{n_1 n_2}{n^2}$ is irrelevant to

the direction of the solution, this proves that OLS (Eq. (3)) is equivalent to the FLDA method (Eq. (2)).

Several works ([16, 36, 17], to cite a few) have extended the above OLS-LDA formulation to multiclass problems. It has been shown that the LDA solution can be obtained from a multivariate regression fit. In particular, Hastie et al. in [16], introduced a richer and more flexible classification scheme to LDA, called *optimal scoring*, which we briefly describe below.

Let $\mathbf{X}$ be as before and $\mathbf{Y}$ be a $n \times K$ matrix of binary variables such that y_{ij} is an indicator variable of whether the i^{th} observation belongs to the j^{th} class. Let us define $\mathbf{\Theta} = [\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_q] \in \mathbb{R}^{K \times q}$, where the vectors $\boldsymbol{\theta}_j$ are recursively obtained, for $j = 1, 2, \ldots, q$, as the solution of the following constrained least squares problem which resumes the optimal scoring:

$$\begin{aligned}(\boldsymbol{\beta}_j, \boldsymbol{\theta}_j) = & \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\theta} \in \mathbb{R}^K}{\arg\min} \|\mathbf{Y}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\beta}\|_2^2, \\ & s.t.\ \frac{1}{n}\boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta} = 1,\ \boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta}_l = 0 \quad \forall l = 1, 2, \ldots, j-1. \end{aligned} \tag{4}$$

Note that when $j = 1$ the orthogonality condition in Eq. (4), which is imposed to avoid trivial solutions, is vacuous and hence it is not enforced. Details about the computational implementation to solve Eq. (4) can be found in [16].

In the sequel we shall refer to $\boldsymbol{\theta}_j$ as the "score vector". Observe that $\boldsymbol{\theta}_j$ is the vector in $\mathbb{R}^K$ for which the mapping from $\mathbb{R}^{n \times K}$ to $\mathbb{R}^n$ defined by $\mathbf{Y} \to \mathbf{Y}\boldsymbol{\theta}_j$, results optimal for the constrained least squares problem defined by (4). This mapping is precisely what introduces more flexibility into the LDA framework since it transforms binary variables into real ones.

Clemmensen et al. [6] introduced a regularized version of the optimal scoring problem by adding two penalization terms to the functional in (4). These penalization terms on one side induce sparsity and on the other side allow correlated variables into the solution. This regularized LDA formulation, named SDA, consists on recursively solving for $j = 1, 2, \ldots, q$, the following problem:

$$\begin{aligned}(\boldsymbol{\beta}_j, \boldsymbol{\theta}_j) = & \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\theta} \in \mathbb{R}^K}{\arg\min} \{\|\mathbf{Y}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda_1\|\boldsymbol{\beta}\|_1 + \lambda_2\|\boldsymbol{\beta}\|_2^2\}, \\ & s.t. \quad \frac{1}{n}\boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta} = 1,\ \boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta}_l = 0 \quad \forall l = 1, 2, \ldots, j-1, \end{aligned} \tag{5}$$

where λ_1 and λ_2 are predefined positive constants, called regularization parameters, which balance the amount of sparsity and the correlation of vari-

ables, respectively. Later on we shall analyse appropriate forms for selecting those regularization parameters.

Problem (5) is alternately and iteratively solved as follows. At first $\boldsymbol{\theta}_j$ is hold fixed and optimization is performed with respect to $\boldsymbol{\beta}_j$. Then $\boldsymbol{\beta}_j$ is hold fixed and optimization is performed with respect to $\boldsymbol{\theta}_j$. The following two steps are iterated:

1. For given (fixed) $\boldsymbol{\theta}_j$, solve:

$$\boldsymbol{\beta}_j = \underset{\boldsymbol{\beta}\in\mathbb{R}^p}{\arg\min}\{\|\mathbf{Y}\boldsymbol{\theta}_j - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda_1\|\boldsymbol{\beta}\|_1 + \lambda_2\|\boldsymbol{\beta}\|_2^2\}. \tag{6}$$

2. For given (fixed) $\boldsymbol{\beta}_j$, solve:

$$\begin{aligned}\boldsymbol{\theta}_j = &\underset{\boldsymbol{\theta}\in\mathbb{R}^K}{\arg\min}\, \|\mathbf{Y}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\beta}_j\|_2^2 \\ &s.t.\ \frac{1}{n}\boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta} = 1,\ \boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta}_l = 0 \quad \forall l = 1, 2, \ldots, j-1.\end{aligned}$$

For computational implementation details of the above steps we refer the reader to [6] and [31].

The solution of problem (5) provides q discriminant direction, $\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, ..., \boldsymbol{\beta}_q$, allowing standard LDA procedure with the $n\times q$ matrix $\left(\boldsymbol{X}\boldsymbol{\beta}_1\ \boldsymbol{X}\boldsymbol{\beta}_2\ \ldots\ \boldsymbol{X}\boldsymbol{\beta}_q\right)$ to be performed.

Solving Eq. (6) involves the well-known elastic-net problem (e-net) [38], which is similar to LASSO (least absolute shrinkage and selection operator) [32], since e-net performs automatic sparse variable selection and it also allows selection of groups of correlated variables (this is due to the introduction of ℓ_2-norm term).

In this work we propose the use of a penalized version of the SDA method to efficiently solve the binary classification problem appearing in BCI systems based on ERPs. This new method seeks to increase classification performance by taking into account information about the difference between classes by means of the inclusion of appropriate anisotropy matrices into the penalizing terms. The use of adaptive penalizers and, in particular of anisotropy matrices in regularization method for inverse ill-posed problems is a new approach that has shown to produce significantly better results than those obtained with the corresponding non-adaptive or isotropic penalizers [26].

3 A new approach: Kullback-Leibler Penalized Sparse Discriminant Analysis

In a pattern recognition problem is very important to analyze the data and be able to extract from them as much prior information as possible since, needless to say, a good classification performance will largely depend on how well the problem is understood through the available data.

Let $\mathbf{X}$, $\mathbf{Y}$, $\boldsymbol{\beta}_j$, $\boldsymbol{\theta}_j$, λ_1 and λ_2 be as before, and let $\mathbf{D}_1$ and $\mathbf{D}_2$ be $p \times p$ diagonal positive definite matrices. The KLSDA scheme consists of recursively solving, for $j = 1, 2, \ldots, q$, the following regularized constrained least squares problem:

$$\begin{aligned}\left(\boldsymbol{\beta}_j, \boldsymbol{\theta}_j\right) = & \underset{\boldsymbol{\beta}\in\mathbb{R}^p, \boldsymbol{\theta}\in\mathbb{R}^K}{\arg\min} \{\|\mathbf{Y}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda_1\|\mathbf{D}_1\boldsymbol{\beta}\|_1 + \lambda_2\|\mathbf{D}_2\boldsymbol{\beta}\|_2^2\}, \\ & s.t. \quad \frac{1}{n}\boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta} = 1,\ \boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta}_l = 0 \quad \forall l = 1, 2, \ldots, j-1. \end{aligned} \tag{7}$$

Here again, for $j = 1$ the orthogonality condition is vacuous. As in the SDA case, the solution of problem (7) is approximated by alternatively iterating the following two steps (with an adequate initialization):

1. Given $\boldsymbol{\theta}_j$ solution of (9), solve:

$$\boldsymbol{\beta}_j = \underset{\boldsymbol{\beta}\in\mathbb{R}^p}{\arg\min}\{\|\mathbf{Y}\boldsymbol{\theta}_j - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda_1\|\mathbf{D}_1\boldsymbol{\beta}\|_1 + \lambda_2\|\mathbf{D}_2\boldsymbol{\beta}\|_2^2\}. \tag{8}$$

2. Given $\boldsymbol{\beta}_j$ solution of (8), solve:

$$\begin{aligned}\boldsymbol{\theta}_j = & \underset{\boldsymbol{\theta}\in\mathbb{R}^K}{\arg\min} \|\mathbf{Y}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\beta}_j\|_2^2 \\ & s.t. \quad \frac{1}{n}\boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta} = 1,\ \boldsymbol{\theta}^T\mathbf{Y}^T\mathbf{Y}\boldsymbol{\theta}_l = 0\ \forall l = 1, 2, \ldots, j-1. \end{aligned} \tag{9}$$

The vector $\boldsymbol{\beta}_j$ solution of (8), not only inherits both the correlated variables selection and sparsity properties of SDA, but it also contains in each one of its components appropriate discriminative information which is suitable for improving separability between classes. As before, the classification rule is constructed based upon the $n \times q$ matrix $\left(\boldsymbol{X}\boldsymbol{\beta}_1\ \boldsymbol{X}\boldsymbol{\beta}_2\ \ldots\ \boldsymbol{X}\boldsymbol{\beta}_q\right)$. In the following subsection we show how the Kullback-Leibler divergence can be used for constructing the anisotropy matrices $\mathbf{D}_1$ and $\mathbf{D}_2$ in such a way that they properly incorporate discriminative information into KLSDA.

3.1 Kullback-Leibler anisotropy matrix

Discriminative information can be incorporated into KLSDA by appropriately quantifying the "distances" between classes, or more precisely, between their probability distributions. Although there is a wide variety of "metrics" for comparing probability distributions [1], we shall use here the well-known Kullback-Leibler divergence [22]. The decision to use this particular "metric" is due not only to its nice mathematical properties, but also to the fact that it was already successfully applied in many classification problems [12, 27, 13].

Let us suppose first that we are dealing with a binary classification problem ($K = 2$). Let $\mathbf{N}$ be a discrete random variable defined on a discrete outcome space $\mathcal{N}$ and consider two probability functions $f_1(n)$ and $f_2(n)$, $n \in \mathcal{N}$. Then, the Kullback-Leibler "distance" (KLD) of f_1 relative to f_2 is defined as:

$$D_{\mathrm{KL}}(f_1||f_2) \doteq \sum_{n\in\mathcal{N}} f_1(n) \log\left(\frac{f_1(n)}{f_2(n)}\right),$$

with the convention that $0.\log 0 \doteq 0$. Although $D_{\mathrm{KL}}(f_1||f_2)$ quantifies the discrepancy between f_1 and f_2, it is not a metric in the rigorous mathematical sense, because it is not symmetric and it does not satisfy the triangle inequality. If, for any reason, symmetry is desired then a modified KLD, called J-divergence, can be defined as follows:

$$J_{\mathrm{KL}}(f_1, f_2) \doteq \frac{D_{KL}(f_1||f_2) + D_{KL}(f_2||f_1)}{2}.$$

For measuring discrepancy between K probability distributions, $f_1, \ldots, f_K$, since KLD is "additive" [22], an appropriate measure can be build up by adding up the J-divergences between all possible pairs of distributions, that is:

$$J_{KL}\left(\{f_j\}_{j=1}^{K}\right) \doteq \sum_{i=1}^{K-1}\sum_{j=i+1}^{K} J_{\mathrm{KL}}(f_i, f_j).$$

Let $f_j^i(\cdot)$ be the probability function of the j^{th} class in the i^{th} sample, with $j = 1, 2, \ldots, K$ and $i = 1, 2, \ldots, p$. We define the J-divergence at sample i as

$$J_{KL}(i) \doteq J_{KL}\left(\left\{f_j^i\right\}_{j=1}^{K}\right). \tag{10}$$

This function quantifies the discrepancy between the K classes at sample i. A value of $J_{KL}(i)$ close to zero means that there is very little discriminative

information at sample i, while a large value of $J_{KL}(i)$ means that sample i contains a significant amount of discriminative information which we definitely want to take into account for constructing the solution vectors $\boldsymbol{\beta}_j$. In Section 5.1 we show in detail how J-divergence is able to highlight the most discriminative samples.

The available a priori discriminative information can be incorporated into the KLSDA formulation (Eq. (7)) by means of appropriately constructed anisotropy matrices $\mathbf{D}_1$ and $\mathbf{D}_2$. Since we wish to stand out those samples containing significant amounts of discriminative information, the matrices $\mathbf{D}_1$ and $\mathbf{D}_2$ must be constructed so as to strongly penalize those samples where there is little or none discriminative information while avoiding penalization at the remaining ones.

3.2 Computational implementation

Our computational implementation of KLSDA bellow is made by appropriately modifying the original SDA algorithm [31]. Thus KLSDA is mainly solved in two steps. In the first step, Eq. (8), which is a generalized version of the e-net problem [28], is solved. The second step consists of updating the optimal score vector $\boldsymbol{\theta}_j$ by solving Eq. (9). It is shown in [6] that the solution of Eq. (9) is given by $\boldsymbol{\theta}_j = s(\mathbf{I} - \boldsymbol{\Theta}_{j-1}\boldsymbol{\Theta}_{j-1}^T\boldsymbol{\pi})\boldsymbol{\pi}^{-1}\mathbf{Y}^T\mathbf{X}\boldsymbol{\beta}_j$, where $\boldsymbol{\Theta}$ is the $K \times q$ matrix containing the score vectors $\boldsymbol{\theta}_j$, $\boldsymbol{\pi} = \frac{1}{n}\mathbf{Y}^T\mathbf{Y}$ and s is a proportionality constant such that $\boldsymbol{\theta}_j\boldsymbol{\pi}\boldsymbol{\theta}_j = 1$.

In regard to the first step, it is known that the e-net problem can be reformulated by means of LASSO. In fact by defining the following augmented variables:

$$\tilde{\mathbf{X}} = \begin{pmatrix} \mathbf{X} \\ \sqrt{\lambda_2\mathbf{D}_2} \end{pmatrix}_{(n+p)\times p}, \quad \tilde{\mathbf{Y}}_j = \begin{pmatrix} \mathbf{Y}\boldsymbol{\theta}_j \\ \mathbf{0}_{p\times 1} \end{pmatrix}_{(n+p)\times 1},$$

the generalized e-net problem (Eq. (8)) can be re-written as:

$$\hat{\boldsymbol{\beta}}_j = \underset{\boldsymbol{\beta}\in\mathbb{R}^p}{\arg\min}\{\|\tilde{\mathbf{Y}}_j - \tilde{\mathbf{X}}\boldsymbol{\beta}\|_2^2 + \lambda_1\|\mathbf{D}_1\boldsymbol{\beta}\|_1\}, \tag{11}$$

which is known as generalized LASSO [33]. If $\mathbf{D}_1$ is invertible the solution of (11) can be found as $\hat{\boldsymbol{\beta}}_j = \mathbf{D}_1^{-1}\hat{\boldsymbol{\alpha}}_j$, where $\hat{\boldsymbol{\alpha}}_j$ is the solution of:

$$\hat{\boldsymbol{\alpha}}_j = \underset{\boldsymbol{\alpha}\in\mathbb{R}^p}{\arg\min}\{\|\tilde{\mathbf{Y}}_j - \tilde{\mathbf{X}}\mathbf{D}_1^{-1}\boldsymbol{\alpha}\|_2^2 + \lambda_1\|\boldsymbol{\alpha}\|_1\}. \tag{12}$$

The LARS-EN algorithm, presented in [38], provides an efficient way of solving problem (12).

3.3 Regularization parameters

It is well-known that in every regularization method the choice of the regularization parameters is crucial. For Tikhonov-type functionals a popular and widely used method for approximating the optimal parameters is the so called L-curve criterion. One of the main advantages of this selection criterion is the fact that it does not require any prior knowledge about the noise. Roughly speaking, the method finds an optimal compromise between the norm of the residual and the norm of the regularized solution by selecting the point of maximal curvature in a log-log plot of those two quantities, parametrized by the regularization parameter. For details see [14].

Despite its popularity, the L-curve method cannot be directly applied to multi-parameter penalization functionals like (8). In 1998, Belge et al. proposed and extension of the L-curve technique for approximating the optimal regularization parameters in those cases, called L-hypersurface technique [2], which we briefly describe next.

Let $\mathbf{X} \in \mathbb{R}^{n\times p}$, $\mathbf{y} \in \mathbb{R}^n$, $\mathbf{R}_i \in \mathbb{R}^{l\times p}$ for $i = 1, \ldots, M$, be all given, and consider the following multi-parameter regularized least squares problem:

$$\boldsymbol{\beta}(\boldsymbol{\lambda}) = \arg\min_{\boldsymbol{\beta}} \{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{i=1}^{M} \lambda_i \|\mathbf{R}_i\boldsymbol{\beta}\|_r^r\}, \tag{13}$$

where $1 \leq r \leq 2$ and let $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_M)$ denote the regularization parameter vector. Define $z(\boldsymbol{\lambda}) \doteq \log \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}(\boldsymbol{\lambda})\|_2^2$ and $x_i(\boldsymbol{\lambda}) \doteq \log \|\mathbf{R}_j\boldsymbol{\beta}(\boldsymbol{\lambda})\|_r^r$, for $i = 1, \ldots, M$. The L-hypersurface associated to problem (13) is the subset $S(\boldsymbol{\lambda})$ of $\mathbb{R}^{M+1}$ parametrized by $\boldsymbol{\lambda} \in \mathbb{R}_+^M$ defined as $S(\boldsymbol{\lambda}) \doteq \{(x_1(\boldsymbol{\lambda}), ..., x_M(\boldsymbol{\lambda}), z(\boldsymbol{\lambda})) : \boldsymbol{\lambda} \in \mathbb{R}_+^M\}$.

The L-hypersurface criterion consists of finding $\boldsymbol{\lambda}^* \in \mathbb{R}_+^M$ such that $S(\boldsymbol{\lambda}^*)$ is the point of maximal Gaussian curvature of the L-hypersurface $S(\boldsymbol{\lambda})$. Although approximating $\boldsymbol{\lambda}^*$ is most of the times very costly from a computational point of view, in [2] the authors show that a good approximation to $\boldsymbol{\lambda}^*$ is given by the minimizer of the residual norm, i.e., by the vector $\hat{\boldsymbol{\lambda}} \in \mathbb{R}_+^M$ satisfying $z(\hat{\boldsymbol{\lambda}}) = \min_{\boldsymbol{\lambda}\in\mathbb{R}_+^M} z(\boldsymbol{\lambda})$.

We describe next the details of how the optimal regularization vector $(\hat{\lambda}_1, \hat{\lambda}_2)$ given by the L-hypersurface approach is approximated within the

KLSDA context formalized in Eq. (7) or, more precisely, in the context of the generalized e-net problem (Eq. (8)). For $j = 1, 2, \ldots, q$, define $z_j(\boldsymbol{\lambda}) = \log \|\mathbf{Y}\boldsymbol{\theta}_j - \mathbf{X}\boldsymbol{\beta}_j(\boldsymbol{\lambda})\|_2^2$, $x_{1,j}(\boldsymbol{\lambda}) = \log \|\mathbf{D}_1\boldsymbol{\beta}_j(\boldsymbol{\lambda})\|_1$ and $x_{2,j}(\boldsymbol{\lambda}) = \log \|\mathbf{D}_2\boldsymbol{\beta}_j(\boldsymbol{\lambda})\|_2^2$. The L-hypersurface associated to (7) is then defined as $S_j(\boldsymbol{\lambda}) = \{(x_{1,j}(\boldsymbol{\lambda}), x_{2,j}(\boldsymbol{\lambda}),\, z_j(\boldsymbol{\lambda})) : \boldsymbol{\lambda} \in \mathbb{R}_+^2\}$.

Although generalized e-net is defined in terms of λ_1 and λ_2, there are other possible choices for tuning parameters [38]. For example, the ℓ_1-norm of the coefficients (t) can be chosen instead of λ_1, In fact, this can be achieved by re-writing the LASSO version (Eq. (12)) of our generalized e-net as a constrained optimization problem with an upper bound on $\|\boldsymbol{\alpha}\|_1$. Similarly, since the LARS-EN algorithm is a forward stagewise additive fitting procedure, the number of steps κ of the algorithm can also be used as a tuning parameter replacing λ_1. This is so because, for each fixed λ_2, LARS-EN produces a finite number of vectors $\boldsymbol{\beta} \in \mathbb{R}^p$ which are approximations of the true solution at each step κ. In our numerical experiment we adopted λ_2 and κ as tuning parameters and, in accordance to Belge's remark described above, the best $\boldsymbol{\beta}_j$ solutions were selected as those minimizing the residual norm. All the steps of the algorithm solving KLSDA with this proposal (together with automatic parameter selection) are presented in Algorithm 1.

4 P300 speller databases

Two real ERP-EEG databases were used to evaluate the classification performance of our KLSDA method.

4.1 Dataset-1

Dataset-1 is an open-access P300 speller database from Laboratorio de Investigación en Neuroimagenología at Universidad Autónoma Metropolitana, Mexico D.F., described in [23]. This database consists of EEG records acquired from 25 healthy subjects, recorded by 10 channels (Fz, C3, Cz, C4, P3, Pz, P4, PO7, PO8, Oz) at 256 Hz sampling rate using the g.tec gUSBamp. During the acquisition, the EEG records were filtered with a Chebyshev Notch 4^{th} order filter with cutoff frequencies of 58-62 Hz and a Chebyshev band-pass 8^{th} order filter with cutoff frequencies of 0.1-60 Hz. A 6-by-6 matrix containing letters and numbers was presented to each subject on a computer screen. During the experiment, the subject was asked to spell dif-

Algorithm 1 KLSDA with automatic parameter selection

Inputs: $\mathbf{X}$, K, q, $\mathbf{Y}$, $\mathbf{D}_1$, $\mathbf{D}_2$, $\Lambda_2 = \left\{\lambda_2^{(1)}, \ldots, \lambda_2^{(d)}\right\}$.

1: Define $\boldsymbol{\pi} = \frac{1}{n}\mathbf{Y}^T\mathbf{Y}$

2: Initialize: $\boldsymbol{\Theta} = eye(K, q)$.

3: **for** $j = 1, \ldots, q$ **do**

4: **while** sparse discriminative direction $\boldsymbol{\beta}_j$ has not converged **do**

5: **for** $i = 1, \ldots, d$ **do**

6: Re-define the variables:

$$\tilde{\mathbf{X}} = \begin{pmatrix} \mathbf{X} \\ \sqrt{\lambda_2^{(i)}\mathbf{D}_2} \end{pmatrix}_{(n+p)\times p}, \quad \tilde{\mathbf{Y}}_j = \begin{pmatrix} \mathbf{Y}\boldsymbol{\theta}_j \\ \mathbf{0}_{p\times 1} \end{pmatrix}_{(n+p)\times 1}$$

7: Solve the generalized e-net problem and save the solution path:

$$(\mathbf{A}_j, \kappa) = \text{LARSEN}(\tilde{\mathbf{X}}_j\mathbf{D}_1^{-1}, \tilde{\mathbf{Y}}_j), \quad \mathbf{B}_j = \mathbf{D}_1^{-1}\mathbf{A}_j$$

8: Find the residual:

$$\mathbf{R}(\lambda_2, 1:\kappa) = \|\mathbf{Y}\boldsymbol{\theta}_j - \mathbf{X}\mathbf{B}_j\|_2^2$$

9: Save the solutions:

$$\mathbf{B}_{\text{all}}(\lambda_2, 1:\kappa, :) = \mathbf{B}_j$$

10: **end for**

11: Select the optimal direction:

$$(\hat{\lambda_2}, \hat{\kappa}) = \arg\min_{\lambda_2,\kappa} \mathbf{R}(\lambda_2, \kappa)$$

$$\boldsymbol{\beta}_j = \mathbf{B}_{\text{all}}(\hat{\lambda_2}, \hat{\kappa}, :)$$

12: Update j^{th} column of $\boldsymbol{\Theta}$: $\boldsymbol{\theta}_j = (\mathbf{I} - \boldsymbol{\Theta}_{(j-1)}\boldsymbol{\Theta}_{(j-1)}^T\boldsymbol{\pi})\boldsymbol{\pi}^{-1}\mathbf{Y}^T\mathbf{X}\boldsymbol{\beta}_j$

13: Normalize to unit length: $\boldsymbol{\theta}_j = \frac{\boldsymbol{\theta}_j}{\sqrt{\boldsymbol{\theta}_j^T\boldsymbol{\pi}\boldsymbol{\theta}_j}}$

14: **end while**

15: **end for**

Outputs: : $[\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_q]$, $\left[\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_q\right]$

ferent words. The person had to focus on one character at the time. As stimulus, a row or a column of the matrix was randomly highlighted for a period of 62.5 ms with inter-stimuli intervals of 125 ms. In each stimulating block (consisting in 12 consecutive flashings), every row and every column of the matrix was intensified only once. For each character to be spelled this stimulating block was repeated 15 times. If the person was well concentrated, a relevant event occurred when the chosen character was illuminated, i.e. an ERP was elicited. Thus, in a binary classification problem, the 6-by-6 matrix generates twelve possible events, of which only two are labelled as target.

Each subject participated in 4 sessions, of which the first two were copy-spelling runs, i.e, they contained the true label data vector, with 21 characters to be spelled. For this reason, in the present work we used those two copy-spelling sessions as our dataset. The EEG records were filtered from 0.1 Hz to 12 Hz by a 4^{th} order forward-backward Butterworth band-pass filter. A 1000 ms data segment was extracted (windowed) from the EEG records at the beginning of each stimulus. A total of 3780 EEG patterns (of which only 630 are target) with dimensionality of $10 \times 256 = 2560$ conforms the dataset per each subject.

4.2 Dataset-2

Dataset-2 is a P300 speller database of patients with amyotrophic lateral sclerosis (ALS) obtained from the Neuroelectrical Imaging and BCI Laboratory, IRCCS Fondazione Santa Lucia, Rome, Italy. The EEG data were recorded using g.MOBILAB and g.Ladybird with active electrode (g.tec, Austria) from eight channels (Fz, Cz, Pz, Oz, P3, P4, PO7 and PO8). All channels were referenced to the right earlobe and grounded to the left mastoid. The EEG signals were digitized at 256 Hz. Eight participants with ALS were required to copy-spell seven predefined words of five characters each using a P300 speller paradigm. As in Dataset-1, a 6-by-6 matrix containing alphanumeric characters was used. The rows and columns of the matrix were randomly intensified for 125 ms, follow by an inter stimulus interval of the same length. In each stimulating block all rows and columns were intensified 10 times. For more details about this dataset we refer the reader to [30].

In a pre-processing stage, the EEG records from each channel were band-pass filtered from 0.1 Hz to 10 Hz using a 4^{th} order Butterworth filter. Then, data segments of 1000 ms were extracted at the beginning of each intensification. The ALS P300 speller dataset consists of 4200 patterns, of which

700 are target, with $256 \times 8 = 2048$ sample points.

5 Experiments and results

In this section we show how the KLSDA method is implemented in the context of the aforementioned real ERP-based BCI classification problem with both datasets described in Section 4. Since we are dealing here with a binary classification problem ($K = 2$), there is only one direction vector $\boldsymbol{\beta}$ in Eq. (7), with which a linear classifier has then to be implemented with the $n \times 1$ projected data vector $\mathbf{X}\boldsymbol{\beta}$.

In the computational implementation of Algorithm 1, the parameter λ_2 was allowed to vary between 10^{-8} and 10^{-1} in a log-scale, and an upper bound of 800 was chosen for the ℓ_1-norm of the coefficients ($t \leq 800$). All codes were run in MATLAB on an AMD FX(tm)-6300 Six-Core PC with 4GB of memory. It is timely to mention here that although the CPU-time needed to find the direction vector was around 2800 secs, once $\boldsymbol{\beta}$ was found, training and testing the linear classifier with the projected data vector $\mathbf{X}\boldsymbol{\beta}$ took only about 3.36 secs.

In experiments we decided to use the symmetric version of KLD, J_{KL}, as a measure of discrepancy. The probability distribution of each class was estimated by using histogram with centered bins between the most minimum and most maximum value in target and non-target training data. We constructed an anisotropy matrix $\mathbf{D} = diag(d_1, \ldots, d_p)$, where $d_i = C/J_{KL}(i),\ i = 1, \ldots, p$, with C being the constant that makes $\det(\mathbf{D}) = 1$, i.e., $C = \left(\prod_{j=1}^{p} J_{KL}(j)\right)^{1/p}$. Note that with this choice, d_i is large where $J_{KL}(i)$ is small and vice-versa. If there exist $i_0, 1 \leq i_0 \leq p$, such that $J_{KL}(i_0) = 0$ then the matrix $\mathbf{D}$ cannot be formally defined as above. This case, however, can be overcomed by simply replacing $J_{KL}(i_0)$ by $J_{KL}(i_0) + \epsilon$, for $i = 1, \ldots, p$, with ϵ very small. The anisotropy matrix $\mathbf{D}$ so defined is clearly diagonal, symmetric and positive definite.

Four different configurations of KLSDA, denoted by KLSDA0, KLSDA1, KLSDA2 and KLSDA3, were implemented. The first one KLSDA0 correspond to the standard SDA approach with automatics parameter selection. The second and third ones, KLSDA1 and KLSDA2, incorporate the anisotropic matrix $\mathbf{D}$ in the ℓ_1 and ℓ_2 penalizers, respectively. Finally, KLSDA3 incorporate the anisotropic matrix $\mathbf{D}$ in both the ℓ_1 and ℓ_2 terms.

Table 1 summarizes these four configurations (where I denotes the $p \times p$ identity matrix). The classification performances of KLSDA method is compared with those obtained with the original SDA approach.

Table 1: KLSDA configurations used.

	KLSDA0	KLSDA1	KLSDA2	KLSDA3
$\mathbf{D_1}$	**I**	**D**	**I**	**D**
$\mathbf{D_2}$	**I**	**I**	**D**	**D**

5.1 Kullback-Leibler divergence for ERP detection

It is reasonable to think that the KLD is an appropriate measure for enhancing the impact of the P300 wave in the KLSDA solution, by selecting both the most discriminative channels and the most discriminative time samples. With this in mind, for both databases, the J-divergence at each sample was estimated as described in Section 3.1 (Eq. (10)). An analysis of this J-divergence as a function of channel and time allowed us to detect which samples were the most discriminative ones. The corresponding plots for six selected subjects of Dataset-1 and six selected subjects of Dataset-2 are presented in Figure 1.

Note that although for most of the subjects the discriminative samples are mostly located in the 250-500 ms window (which in fact corresponds to the latency window of the P300 wave), in other cases (e.g. subject 7 and 14 from Dataset-1 and subject 1 from Dataset-2) the most discriminative samples are somewhat randomly distributed all over the plot. Moreover, there are cases in which for some channels the J-divergence shows no contribution at all to class separation (e.g. subject 8 from Dataset-1 and subject 2 from Dataset-2). Figure 1 also shows the high variability of the ERP morphology between subjects as it was pointed out in Section 1.

5.2 Classification results

There are different available measures for evaluating a BCI classification method [3]. The receiver operator characteristics (ROC) curve is a powerful tool for evaluating a two-class unbalanced problem [9]. In the present

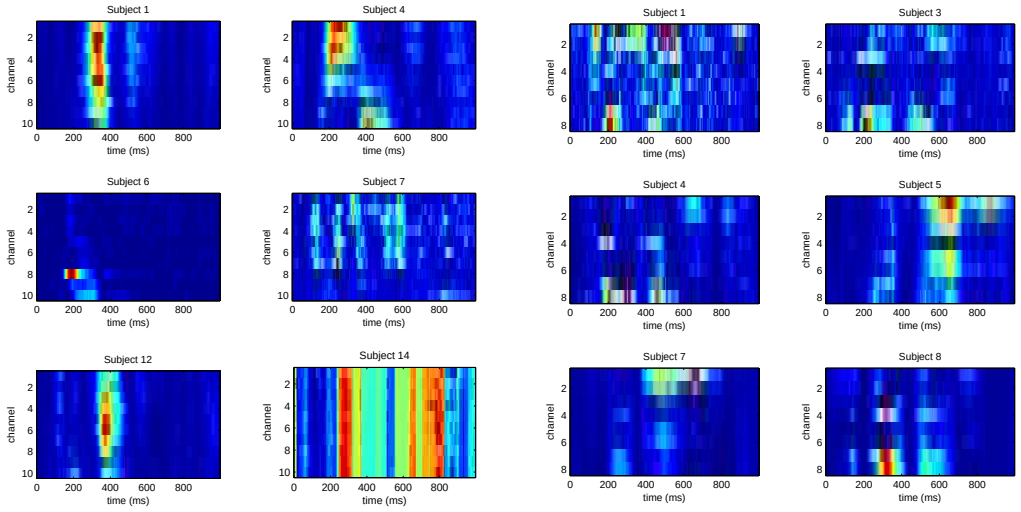

(a) Dataset-1 (b) Dataset-2

Figure 1: Discriminative information by means of J-divergence at each sample point in the channel-time plane for six selected subjects from Dataset-1 and Dataset-2.

work the Area Under the ROC Curve [5], denoted by AUC, was used as the classification performance measure. For avoiding classification bias, a 3-fold cross-validation procedure was implemented.

Figure 2 shows the classification results obtained for each subject from Dataset-1 with the KLSDA0, KLSDA1, KLSDA2 and KLSDA3 methods. Several remarks are in order. First of all, the method which results in the best classification performance seems to be subject dependent. A second observation is that the poorest classification performance corresponds precisely to those subjects with "randomly spread J-divergence". A final observation is that in nineteen of the twenty five cases there is at least one no trivial KLSDA configuration that outperforms the pure SDA (KLSDA0), while in the remaining six cases the performances are essentially equal.

A similar analysis can be made on the results obtained with Dataset-2 as depicted in Figure 3. Note that in this dataset, in only three cases one KLSDA configuration outperforms SDA.

The average classification results for each KLSDA configuration are presented in the last column of Figure 2 and Figure 3 for Dataset-1 and Dataset-2, respectively. It is worth noting that for Dataset-1 the average classification result were over 77% and for Dataset-2 were over 75% for all KLSDA configuration. These results are very encouraging since an efficient BCI system requires of an accuracy above 70% to allow communication and device control [24].

We have also used both datasets to test the FLDA classifier with 3-fold cross-validation. As expected, due to the curse-of-dimensionality, the classification results were very poor. In fact, for Dataset-1 the average classification performance was only around 60% while for Dataset-2 it was near 65%. These results clearly indicate that regularization improves classification performance.

5.3 Sparsity analysis

Given that solution sparsity was desired, the mean of the number of non-zeros values for each KLSDA configuration was analysed. For the KLSDA0, KLSDA1, KLSDA2 and KLSDA3 methods, the mean of the percentage of non-zero values respect to the number of sample points for Dataset-1 were found to be 5.75%, 7.48%, 5.78% and 7.79%, respectively, while for Dataset-2 those values were found to be 12.27%, 13.03%, 12.98% and 13.19%. Note that for Dataset-2 sparsity is consistently and significantly lower. This fact most

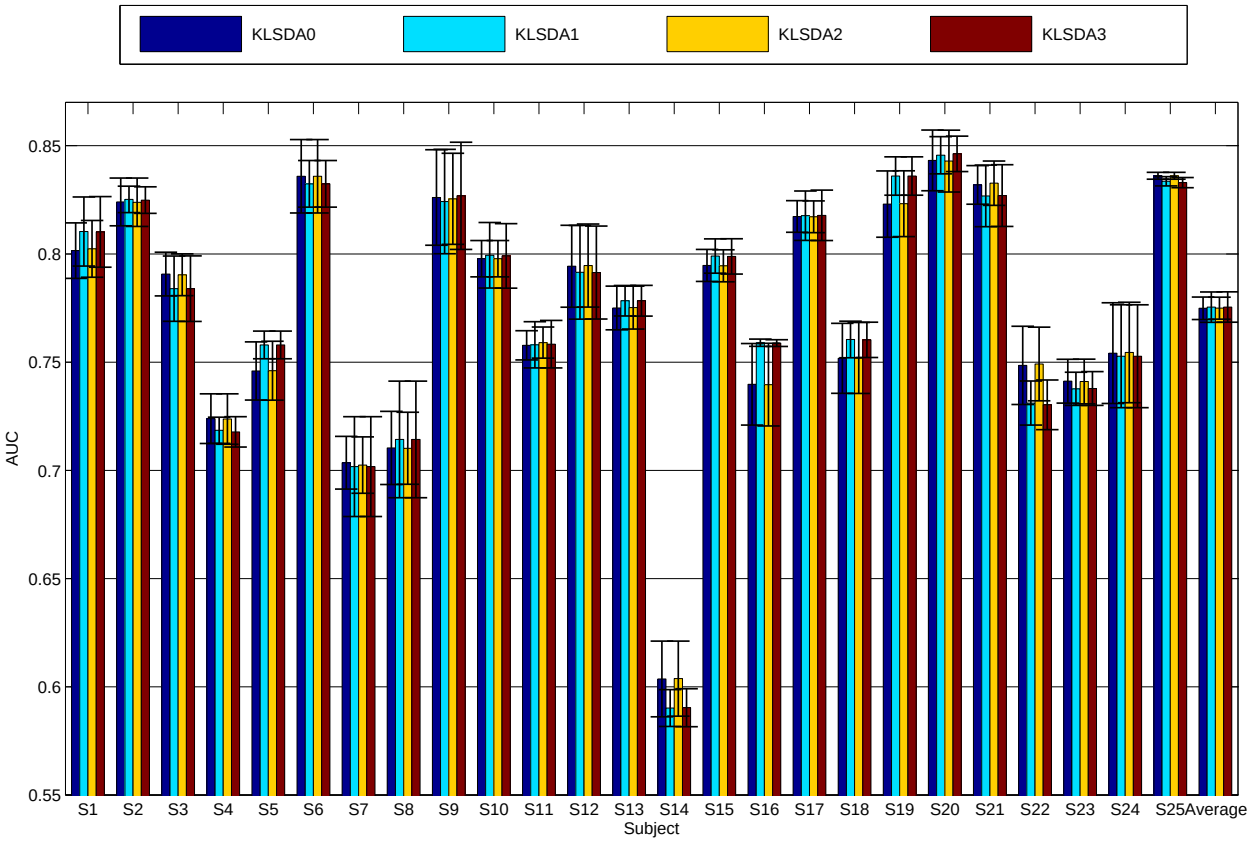

Figure 2: Area under the ROC curve (AUC) on test data from Dataset-1, derived by KLSDA0, KLSDA1, KLSDA2, KLSDA3 and evaluated by 3-fold cross-validation. The errorbar of each subject denotes the standard deviation of AUC on the 3-fold. The errorbar of the average denotes the standard deviation on all subject.

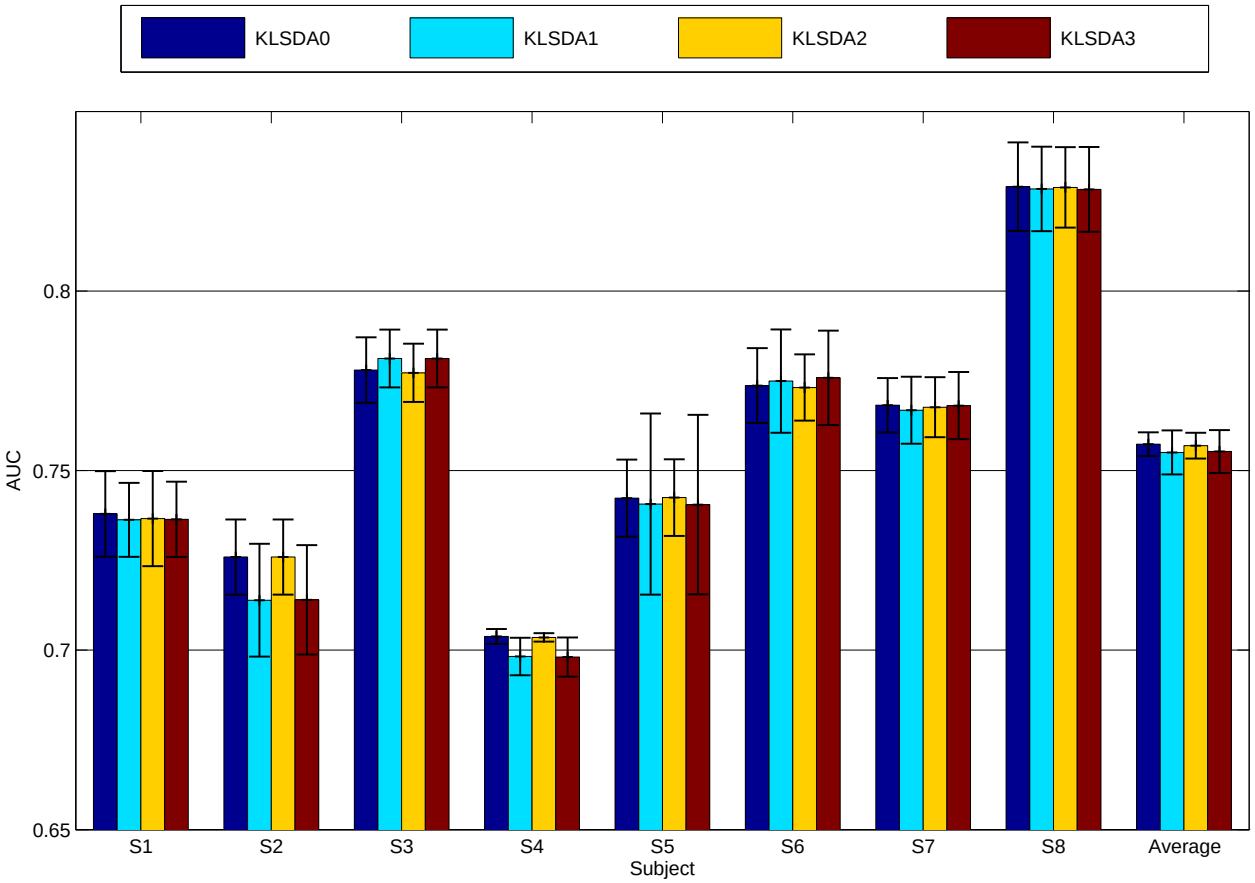

Figure 3: Area under the ROC curve (AUC) on test data from each subject of Dataset-2 in each KLSDA configuration. The errorbar of each subject denotes the standard deviation of AUC on the 3-fold. The errorbar of the average denotes the standard deviation on all subject.

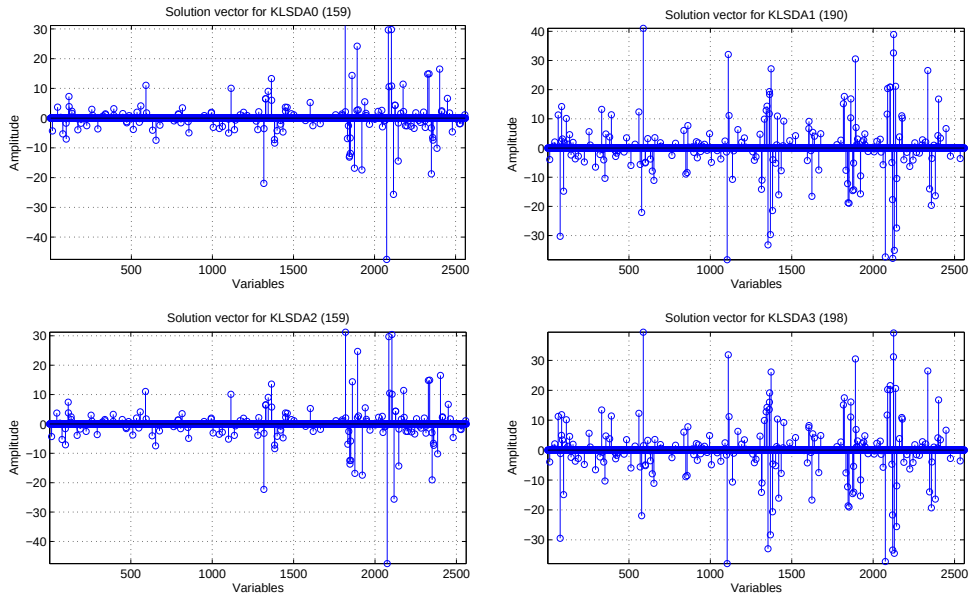

Figure 4: Solution vectors for one subject of Dataset-1 in each KLSDA configuration.

probably reflects the fact that ALS patients database (non-healthy subjects) involved more complex patterns. Finally, it is highly remarkable that all KLSDA configurations with such low percentages of non-zero values achieved such satisfactory classification performances. This observation allows us to conclude that KLSDA constitutes a very robust classifier method.

Figure 4 depicts the morphology of the solution vectors for one subject of Dataset-1 in the four different KLSDA configurations. Between parentheses the number of non-zero values in the solution vector is shown. Note that this number is higher when the anisotropy matrix $\mathbf{D}$ is used in the ℓ_1 penalization term (KLSDA1 and KLSDA3), and also the amplitudes of the coefficients increase in those two cases.

6 Discussion

It is well known that LDA is a commonly used method for ERP classification purposes. Its wide use is mainly due to its robustness, simplicity and

good classification performances. When the number of observations is significantly lower than their dimensionality, LDA performs poorly, reason for which several alternatives to this method have been proposed. In the present work we described different LDA approaches from the statistical literature and developed a new penalized sparse discriminant analysis method called Kullback-Leibler Penalized Sparse Discriminant Analysis. This new method not only inherits the good properties of SDA, but it also allows us to incorporate KLD-based discriminative information, in order to enhance classification performance.

It is important to highlight that our implementation of KLSDA incorporates automatic tuning parameter selection. In light of the sparsity degree of the solution and the classification performances obtained, this procedure has proved to be very adequate for choosing the two regularization parameters in the model.

We tested the new KLSDA approach with two real ERP-EEG datasets. An analysis of the classification results indicates that the KLSDA configuration leading to the best performance is subject depended. In particular, the classification results from Dataset-2 (ALS patients), show that adding KLD information into the solution may not necessarily result beneficial, specially when it provides no clear discriminative information. In regard to this, we shall compare these results with the ones obtained by using another discriminant information measure, as the asymmetric KLD or the ones presented in [20, 29], to cite a few.

It is remarkable that in those cases where KLSDA outperformed SDA there were no computational cost added. In fact, in our experiments we found that SDA's computational cost was the same or even higher than KLSDA's with two anisotropic penalizing terms (KLSDA3).

The results achieved by applying KLSDA in the context of ERP classification problems are high enough (over 75%) to ensure good communication between the brain of the person and the device being controlled. These results encourage us to continue research efforts. There is clearly much room for improvement. Further research is currently underway in several directions. For instance, different discrepancy measures, anisotropy matrices and penalizing terms can be considered.

A final remark is that, although the KLSDA method was inspired by the idea of solving the binary ERP classification problem in BCI systems, it can clearly be applied to any type of classification problem.

Acknowledgements

This work was supported in part by Consejo Nacional de Investigaciones Científicas y Técnicas, CONICET, through PIP 2014-2016 No. 11220130100216-CO, the Air Force Office of Scientific Research, AFOSR/SOARD, through Grant FA9550-14-1-0130 and and by Universidad Nacional del Litoral, UNL, through CAID-UNL 2011 Project No.525 within PACT "Señales, Sistemas e Inteligencia Computacional".

References

References

[1] M. Basseville. Distance measures for signal processing and pattern recognition. *Signal Processing*, 18:349–369, 1989.

[2] Murat Belge, Misha E Kilmer, and Eric L Miller. Simultaneous multiple regularization parameter selection by means of the L-hypersurface with applications to linear inverse problems posed in the wavelet transform domain. In *SPIE's International Symposium on Optical Science, Engineering, and Instrumentation*, pages 328–336. International Society for Optics and Photonics, 1998.

[3] Martin Billinger, Ian Daly, Vera Kaiser, Jing Jin, Brendan Z Allison, Gernot R Müller-Putz, and Clemens Brunner. Is it significant? guidelines for reporting BCI performance. In *Towards Practical Brain-Computer Interfaces*, pages 333–354. Springer, 2012.

[4] Benjamin Blankertz, Steven Lemm, Matthias Treder, Stefan Hauf, and Klaus Robert M*ü*ler. Single-trial analysis and classification of ERP component- a tutorial. *Neuroimage*, 56:814–825, 2011.

[5] Andrew P Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.

[6] Line Clemmensen, Trevor Hastie, Daniela Witten, and Bjarne Ersbøll. Sparse discriminant analysis. *Technometrics*, 53:406–413, 2012.

[7] Richard O Duda, Peter E Hart, and David G Stork. *Pattern classification*. John Wiley & Sons, 2012.

[8] Lawrence Ashley Farwell and Emanuel Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalography and clinical Neurophysiology*, 70(6):510–523, 1988.

[9] Tom Fawcett. An introduction to ROC analysis. *Pattern recognition letters*, 27(8):861–874, 2006.

[10] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936.

[11] Ivan E Gareis, Ruben C Acevedo, Yanina V Atum, Gerardo G Gentiletti, Veronica Medina Banuelos, and Hugo L Rufiner. Determination of an optimal training strategy for a BCI classification task with LDA. In *Neural Engineering (NER), 2011 5th International IEEE/EMBS Conference on*, pages 286–289. IEEE, 2011.

[12] Will Gersch, F Martinelli, J Yonemoto, MD Low, and JA Mc Ewan. Automatic classification of electroencephalograms: Kullback-leibler nearest neighbor rules. *Science*, 205(4402):193–195, 1979.

[13] Anjum Gupta, Shibin Parameswaran, and Cheng-Han Lee. Classification of electroencephalography (EEG) signals for different mental activities using kullback leibler (kl) divergence. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, pages 1697–1700. IEEE, 2009.

[14] Per Christian Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM review*, 34(4):561–580, 1992.

[15] Trevor Hastie, Andreas Buja, and Robert Tibshirani. Penalized discriminant analysis. *The Annals of Statistics*, pages 73–102, 1995.

[16] Trevor Hastie, Robert Tibshirani, and Andreas Buja. Flexible discriminant analysis by optimal scoring. *Journal of the American statistical association*, 89(428):1255–1270, 1994.

[17] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer Series in Statistics, 2009.

[18] Steven A Hillyard and Marta Kutas. Electrophysiology of cognitive processing. *Annual review of psychology*, 34(1):33–61, 1983.

[19] Johannes Hohne, Benjamin Blankertz, Klaus-Robert Muller, and Daniel Bartz. Mean shrinkage improves the classification of ERP signals by exploiting additional label information. In *Pattern Recognition in Neuroimaging, 2014 International Workshop on*, pages 1–4. IEEE, 2014.

[20] Don Johnson and Sinan Sinanovic. Symmetrizing the kullback-leibler distance. *IEEE Transactions on Information Theory*, 2001.

[21] Dean J Krusienski, Eric W Sellers, François Cabestaing, Sabri Bayoudh, Dennis J McFarland, Theresa M Vaughan, and Jonathan R Wolpaw. A comparison of classification techniques for the P300 speller. *Journal of neural engineering*, 3(4):299, 2006.

[22] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[23] Claudia Ledesma-Ramirez, Erik Bojorges-Valdez, Oscar Yáñez-Suarez, Carolina Saavedra, Laurent Bougrain, and Gerardo Gabriel Gentiletti. An open-access P300 speller database. In *Fourth International Brain-Computer Interface Meeting*, 2010.

[24] K Li, V N Raju, R Sankar, Y Arbel, and E Donchin. Advances and challenges in signal analysis for single trial P300-BCI. In *Foundations of Augmented Cognition. Directing the Future of Adaptive Systems*, volume 2, pages 87–94, ll, 2011.

[25] Fabien Lotte, Marco Congedo, Anatole Lécuyer, Fabrice Lamarche, and Bruno Arnaldi. A review of classification algorithms for EEG-based brain-computer interfaces. *Journal of neural engineering*, 4(2):R1, 2007.

[26] Gisela L Mazzieri, Ruben D Spies, and Karina G Temperini. Mixed spatially varying L2-BV regularization of inverse ill-posed problems. *Journal of Inverse and Ill-posed Problems*, 23(6):571–585, 2015.

[27] Pedro J Moreno, Purdy P Ho, and Nuno Vasconcelos. A kullback-leibler divergence based kernel for SVM classification in multimedia applications. In *Advances in neural information processing systems*, page None, 2003.

[28] Geoffroy Mouret, Jean-Jules Brault, and Vahid Partovinia. Generalized elastic net regression. 2013.

[29] Mark D Reid and Robert C Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12(Mar):731–817, 2011.

[30] Angela Riccio, Luca Simione, Francesca Schettini, Alessia Pizzimenti, Maurizio Inghilleri, Marta Olivetti Belardinelli, Donatella Mattia, and Febo Cincotti. Attention and P300-based BCI performance in people with amyotrophic lateral sclerosis. *Front Hum Neurosci 2013 Nov; 7*, 732(2,906), 2013.

[31] Karl Sjöstrand, Line Harder Clemmensen, Rasmus Larsen, and Bjarne Ersbøll. SpaSM: A matlab toolbox for sparse statistical modeling. *Journal of Statistical Software Accepted for publication*, 2012.

[32] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[33] Ryan Joseph Tibshirani and Jonathan E Taylor. The solution path of the generalized lasso. *The Annals of Statistics*, 39:1335–1371, 2011.

[34] Jonathan Wolpaw and Elizabeth Winter Wolpaw. *Brain-Computer Interfaces: principles and practice*. Oxford University Press, USA, 2012.

[35] Jonathan R Wolpaw, Niels Birbaumer, Dennis J McFarland, Gert Pfurtscheller, and Theresa M Vaughan. Brain-computer interfaces for communication and control. *Clinical neurophysiology*, 113(6):767–791, 2002.

[36] Jieping Ye. Leaste squares linear discriminant analysis. In *Proceedings of the 24th international conference on Machine learning*, pages 1087–1093. ACM, 2007.

[37] Yu Zhang, Guoxu Zhou, Jing Jin, Qibin Zhao, Xingyu Wang, and Andrzej Cichocki. Aggregation of sparse linear discriminant analyses for event-related potential classification in brain-computer interface. *International journal of neural systems*, 24(01):1450003, 2014.

[38] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.